

How Much of the Web Is Archived?

Scott G. Ainsworth, Ahmed AlSum, Hany SalahEldeen,
Michele C. Weigle, Michael L. Nelson
Old Dominion University Norfolk, VA, USA
{sainswor, aalsum, hany, mweigle, mln}@cs.odu.edu

ABSTRACT

Although the Internet Archive's Wayback Machine is the largest and most well-known web archive, there have been a number of public web archives that have emerged in the last several years. With varying resources, audiences and collection development policies, these archives have varying levels of overlap with each other. While individual archives can be measured in terms of number of URIs, number of copies per URI, and intersection with other archives, to date there has been no answer to the question "How much of the Web is archived?" We study the question by approximating the Web using sample URIs from DMOZ, Delicious, Bitly, and search engine indexes; and, counting the number of copies of the sample URIs exist in various public web archives. Each sample set provides its own bias. The results from our sample sets indicate that range from 35%-90% of the Web has at least one archived copy, 17%-49% has between 2-5 copies, 1%-8% has 6-10 copies, and 8%-63% has more than 10 copies in public web archives. The number of URI copies varies as a function of time, but no more than 31.3% of URIs are archived more than once per month.

Categories and Subject Descriptors

H.3.7 [Information Storage and Retrieval]: Digital Libraries

General Terms

Design, Experimentation, Standardization

Keywords

Web Architecture, HTTP, Resource Versioning, Web Archiving, Temporal Applications, Digital Preservation

1. INTRODUCTION

With more and more of our business, academic, and cultural discourse contained primarily or exclusively on the web, there has been an increased attention to the problem of web archiving. The focal point of such discussions is the Internet Archive's Wayback Machine, which began archiving the web in 1996 and as of 2010 had over 1.5 billion unique URIs [17], thus making it the largest, longest-running and most well known of publicly available web archives. On the other hand, given the increased attention to the problem of digital preservation, the proliferation of production quality web crawling tools, and the falling cost of resources required for preservation (e.g., storage and bandwidth), there has been a proliferation of additional public web archives

at universities, national libraries, and other organizations. They differ in a variety of ways, including scale, ingest models, collection development policies, software employed for crawling and archiving.

Anecdotally we know that these additional archives have some degree of overlap with the Internet Archive's Wayback Machine, but at the same time they are also not proper subsets of the Wayback Machine. This leads to a common question often asked of those working in digital preservation: "How much of the Web is archived?" We are unaware of any prior attempt to address this question. We do know the approximate size of the Wayback Machine, and we do know that Google has measured the web to contain at least one trillion unique URIs (though they do not claim to index all of them)¹.

In this paper, we do not attempt to measure the absolute size of various public web archives. Instead, we sample URIs from a variety of sources (DMOZ, Delicious, Bitly, and search engine indices) and report on the number of URIs that are archived, the number and frequency of different timestamped versions of each archived URI, and the overlap between the archives. From this, we extrapolate the percentage of the surface web that is archived.

2. RELATED WORK

Although the need for Web archiving has been understood since nearly the dawn of the Web [3], these efforts are for the most part independent in motivation, requirements, and scope. The Internet Archive², the first archive to attempt global scope, came into existence in 1995 [10]. Since then, many other archives have come into existence.

Recognizing the need for coordination, McCown and Nelson proposed the Web Repository Description Framework and API [14]. Still, there is much to be accomplished before the state of Web archiving is understood and a coordinated effort.

Although much has been written on the technical, social, legal, and political issues of Web archiving, little research has been conducted on the archive coverage provided by the existing archives. Day [4] surveyed a large number of archives while investigating the methods and issues associated with

¹<http://googleblog.blogspot.com/2008/07/we-knew-web-was-big.html>

²<http://www.archive.org>

archiving. Day however does not address coverage. Thelwall touches on coverage when he addresses international bias in the Internet Archive [19], but does not directly address the percent of the Web that is covered. McCown and Nelson do address coverage [13], but their research is limited to search engines caches.

Another aspect of coverage still in open discussion is determining what should be covered. When Gomes, Freitas, et al. addressed the design of a national Web archive [6], incompleteness was inherent in their compromise design. Mason argues the Web and digital culture has changed our sense of permanence, which in turn has changed the collecting practices the National Library of New Zealand. [11]. Phillips systematically address the question “what should be archived?”, but concludes that consensus has not been reached and that the Web’s huge volume puts complete archiving out of reach.

Another aspect of Web archives that remained resolved until recently was lack of a standard API. Van de Sompel, et al. have addressed this with Memento [21, 22]. Memento is an HTTP-based framework that bridges Web archives with current resources; it provides a standard API for identifying and dereferencing archived resources through datetime negotiation. In Memento, each original resource (identified with “URI-R”) has 0 or more archived representations (identified with URI-M_{*i*} (*i* = 1..*n*)) that encapsulate the URI-R’s state at times *t_i*. Using the Memento API, clients are able to request URI-M_{*i*} for a specified URI-R. Memento is now an IETF Internet Draft [20].

3. EXPERIMENT

From late November 2010 through early January 2011, we performed an experiment to estimate the percentage of URIs that are represented in Web archives. The primary purpose of this experiment was to estimate the percentage of all publicly-visible URIs that have archive copies available in public archives such as the Internet Archive. A secondary purpose was to evaluate the quality of Web archiving. This experiment was accomplished in three parts:

- Selecting a sample set of URIs that are representative of the Web as a whole,
- Determining the current state of the sample URIs (URI http status, SE index, and estimated age), and
- Discovering Mementos for the sample URIs.

3.1 Sample URI Selection

Discovering every URI that exists is impossible; representative sampling is required. In previous research several methods were used to select a sample of URIs that are representative of the Web as a whole. This experiment’s sampling followed several of these approaches. Like many other research projects, we used the Open Directory Project (DMOZ). We also sampled from search engines using the code and procedures provided by Bar-Yosef [2] and sampled the Delicious Recent bookmark list and Bitly. The reasoning behind these sources and the methods used for URI selection are detailed below.

For practical reasons (e.g., search engine query limits and execution time) we selected a sample size of 1,000 URIs for

Table 1: The four sample sets, all with $n = 1000$

| Collection | Mean | SD | SE |
|---------------------|-------|--------|-------|
| DMOZ (URI-M>0) | 62.68 | 123.86 | 7.68 |
| DMOZ (all) | 56.85 | 119.35 | 7.40 |
| Delicious (URI-M>0) | 81.44 | 232.02 | 14.38 |
| Delicious (all) | 79.40 | 229.45 | 14.38 |
| Bitly (URI-M>0) | 41.64 | 229.18 | 14.20 |
| Bitly (all) | 14.66 | 137.30 | 14.20 |
| SE (URI-M>0) | 6.99 | 25.40 | 1.57 |
| SE (all) | 5.40 | 22.55 | 1.40 |

each of the four sources. Table 1 shows the mean number of mementos for each of the 1000 URI-Rs (both with and without URI-Rs with zero mementos) in each sample set, the standard deviation, and the standard error at a 95% confidence level.

3.1.1 Open Directory Project (DMOZ) Sampling

Using the Open Directory Project (DMOZ)³ as a URI sample source has a long history [16, 7, 18, 7]. Although it is an imperfect source for many reasons (e.g. its contents appear to be driven by commercial motives and are likely biased in favor of commercial sites), DMOZ was included because it provides for comparability with previous studies and because it is one of the oldest sources available. In particular, DMOZ archives dating back to 2000 are readily available, which makes DMOZ a reliable source for old URIs that may no longer exist.

Our URI selection from DMOZ differs from previous methods such as Gulli and Signorini [7] in that we used the entire available DMOZ history instead of a single snapshot in time. In particular, we extracted URIs from every DMOZ archive available⁴, which includes 100 snapshots of DMOZ made from July 20, 2000 through October 3, 2010. First, a combined list of all unique URIs was produced by merging the 100 archives. During this process, the date of the DMOZ archive in which each URI first existed was captured. This date is later used as indirect evidence of the URI’s creation date. From this combined list, 3,806 invalid URIs (URIs not in compliance with RFC 3986) were excluded, 1,807 non-HTTP URIs were excluded, and 17,681 URIs with character set encoding errors were excluded. This resulted in 9,415,486 unique, valid URIs from which to sample. Note that URIs in the DMOZ sample are biased because Internet Archive uses the DMOZ directory as a seed for site-crawling [8].

3.1.2 Delicious Sampling

The next source for URIs is social bookmarking site. In this paper, sampling from Delicious⁵ was used. Delicious is a social bookmarking service started in 2003; it allows users to tag, save, manage and share web pages from a centralized source. Delicious provides two main types of bookmarks. Delicious recent bookmarks are the URIs that have been recently added by users. Delicious popular bookmarks are the currently most popular bookmarks in the Delicious bookmarks set. In our experiment, we retrieved 1,000 URIs from

³<http://www.dmoz.org>

⁴<http://rdf.dmoz.org/rdf/archive>

⁵<http://www.delicious.com>

the Delicious Recent Random URI Generator⁶ (as of Nov. 22, 2010). We also considered the Delicious Popular Random URI Generator. However, it’s small set of distinct URIs which didn’t provide a good sample.

3.1.3 Bitly Sampling

The Bitly⁷ project is a web-based service for URI shortening. Its popularity grew as a result of being the default URI shortening service on the microblogging service Twitter (from 2009-2010), and now enjoys a significant user base of its own. Any link posted on Twitter is automatically shortened and reposted. Bitly creates a “short” URI that when dereferenced issues an HTTP 301 redirect to a target URI. The shortened URI consists of a hash value of up to six alphanumeric characters appended to `http://bit.ly/`, for example the hash value A produces:

```
% curl -I http://bit.ly/A
HTTP/1.1 301 Moved
Date: Sun, 30 Jan 2011 16:00:48 GMT
Server: nginx
Location: http://www.wieistmeineip.de/ip-address
...
```

The shortened URI consumes fewer characters in an SMS message (e.g., a Tweet), protects long URIs with arguments and encodings from being mangled in emails and other contexts, and provides an entry point for tracking clicks by appending a “+” to the URI: `http://bit.ly/A+`. This tracking page reveals when the short URI was created, as well as the dereferences and associated contexts for the dereferences. The creation time of the Bitly is assumed to be greater than or equal to the creation time of the target URI to which the Bitly redirects.

To sample Bitly, we randomly created a series candidate hash values, dereferenced the corresponding Bitly URI, and recorded the target URIs (i.e., the URI in the *Location:* response header). The first 1000 bitlys that returned HTTP 301 responses were used. We also recorded the creation time of the Bitlys via their associated “+” pages.

3.1.4 Search Engine Sampling

Search engines play an important role in web page discovery for most casual users of the Web. Previous studies have examined the relationship between the Web as a whole and the portion indexed by search engines. A search engine sample should be an excellent representation of the Web as a whole. However, the sample must be random, representative, and unbiased. One way to tackle the randomness of this sample is by providing the search engines with multiple random queries, getting the results and choosing again at random from them. This intuitive approach is feasible but suffers from several deficiencies and is extremely biased. The deficiencies reside in the necessity of creating a completely diverse query list of all topics and keywords. Also search engines are normally limited to providing only about the first 1,000 results. Bias, on other hand, comes from the fact

that search engines present results with preference to their page rank. The higher the popularity of a certain URI, and its adherence to the query, the more probable it will appear first in the returned results.

It is necessary to sample the search engine index efficiently, at random, covering most aspects, while also removing ranking bias and popularity completely. Several studies have investigated solving different aspects of this problem. The most suitable solution was presented by Bar-Yossef and Gurevich [2].

As illustrated in the Bar-Yossef paper, there are two methods to implement this unbiased random URL sampler from search engine’s index. The first is by utilizing a pool of phrases to assemble queries that will be later fed to the search engine. The other approach is based on random walks and does not need a preparation step. The first method was utilized in this paper with a small modification to the first phase of pool preparation.

Originally a huge corpus should be assembled, from which the query pool will be created. Instead, the N-grams query list [5] obtained from Google was utilized, and a size of 5 grams were chosen. A total query pool of 1,176,470,663 queries was collected. A random sampling of the queries was provided to the URI sampler as the second phase. A huge number of URIs were produced, and 1,000 were filtered at random to be utilized as the unbiased, random and representative sample of the indexed web.

3.2 Current URI State Determination

After selecting the different sample sets of URIs, the current state of each URI on the web was determined using four components: the current existence (or non-existence) of the resource and the current index of the URI in three search engines: Google, Bing, and Yahoo!.

First, the current status of each URI was tested using the `curl` command. The status was classified into one of six categories based on the HTTP response code. It was important to divide the 3xx responses into two categories because a search engine could carry the redirect information and report the original URI as indexed and display the new URI. Table 2 lists the results of the current status of the URIs on the web. 95% of Delicious and search engine sample URIs sets are live URIs (status code 200), but only 50% of DMOZ and Bitly sample URIs sets are live. The reason is that the Delicious sample came from the “Recent” added bookmarks which means these URIs are still alive and some users are still tagging them. Search engines purge their caches soon after they discover URIs that return a 404 response [15].

Second, we tested the status of each sample URI to see if it was indexed by each search engine. For both Bing and Yahoo search engines, Bing APIs⁸ and Yahoo BOSS APIs⁹, respectively were used to find the URI in these search engine indexes. Note that the APIs return different results than those from the web interfaces [12]. Google indexing

⁶<http://www.delicious.com/recent/?random>

⁷<http://bit.ly>

⁸<http://www.bing.com/developers>

⁹<http://developer.yahoo.com/search/boss/>

Table 2: Sample URIs status on the current Web.

| Status | DMOZ | Delicious | Bitly | SE |
|-----------|------|-----------|-------|-----|
| 200 | 507 | 958 | 488 | 943 |
| 3xx⇒200 | 192 | 27 | 243 | 17 |
| 3xx⇒Other | 50 | 1 | 36 | 3 |
| 4xx | 135 | 8 | 197 | 16 |
| 5xx | 4 | 3 | 6 | 0 |
| Timeout | 112 | 3 | 30 | 21 |

Table 4: Sample URIs status on the search engines.

| | DMOZ | Delicious | Bitly | SE |
|--------------------|------|-----------|-------|-----|
| Bing | 495 | 953 | 218 | 552 |
| Yahoo | 410 | 862 | 225 | 979 |
| Google(APIs Only) | 307 | 883 | 243 | 702 |
| Google(APIs+Cache) | 545 | 951 | 305 | 732 |

task depends on Google Research Search Program APIs¹⁰. In this experiment, a significant difference in the indexed results between Google web interface and Google Research Search Program APIs was found. Table 4 lists the number of the URIs indexed by the different search engines. Google indexed status is shown on two rows; the first one is the number discovered by the API, and the second one is the union between the URIs discovered by the APIs and the URIs found by the Memento proxy on the Google cache.

3.3 Memento Discovery

Concurrently with determining current status, memento discovery was conducted for each URI in the sample sets. In the memento discovery task, the Memento project’s archive proxies and aggregator [21] were used to get a list of all the accessible, public archived versions (mementos) for each URI. Memento’s¹¹ approach is based on a straightforward extension of “Hypertext Transfer Protocol (HTTP)” that results in a way to seamlessly navigate current and prior versions of web resources which might be held by web archives. For any URI, Memento’s *TimeBundle* provides an aggregation of all mementos available from an archive. The ODU Computer Science Memento Aggregator implementation was used to retrieve URI TimeBundles, which searched a large number of archives through archive-specific proxies. For each URI, the proxy queries the archive for the available Mementos for the URI, and returns to the aggregator a list of mementos. The aggregator merges the results from all proxies, then sorts it by date and returns a Timemap for the URI. Table 5 describes a list of the archives that used in this task.

3.4 URI Age Estimation

Intuition is that the longer a URI has been available on the Web, the mementos it will have. Unfortunately, the actual creation date of a URI is almost always unavailable¹². So to estimate the age of the URI, we estimate the creation date with the earliest value of: the date of the first memento, the

¹⁰<http://www.google.com/research/university/search/>

¹¹<http://www.mementoweb.org/>

¹²For a discussion, see: <http://ws-dl.blogspot.com/2010/11/2010-11-05-memento-datetime-is-not-last.html>

Table 3: Number of mementos per URI.

| Mementos per URI | DMOZ | Delicious | Bitly | SE |
|------------------|------|-----------|-------|-----|
| 0 (Not archived) | 93 | 25 | 648 | 225 |
| 1 | 46 | 79 | 100 | 336 |
| 2 to 5 | 142 | 491 | 171 | 320 |
| 6 to 10 | 85 | 35 | 17 | 35 |
| More than 10 | 634 | 370 | 64 | 84 |

date of first DMOZ archive that reported the URI, and the first date the URI was added to Delicious.

4. RESULTS

The results in this section will be presented by sample set. Figure 1 is a histogram of the number of memento URIs retrieved by the Memento proxy servers. This figure shows the distribution of the mementos through time. The DMOZ and Delicious samples have a similar distribution. This figure also shows that there is a low coverage for the period before 2000, as could be expected. The end of 2010 has a good coverage led by the search engine caches.

Figure 2 is a detailed graph of the distribution of the mementos over the time. Each dot represents a single memento. Dot color indicates the source of the memento: Internet Archive, search engine caches, and other archives. Reach row on the y-axis represents a single URI-R, which are ordered bottom-to-top by estimated creation date. The x-axis is the memento’s datetime. Figure 2 shows that most of the mementos before 2008 are provided by the Internet Archive. Search engine caches provide very recent copies of the URI.

Table 3 shows the number of mementos per URI. The Bitly sample has many URIs that are not covered at all. This means that Bitly URIs have not been discovered by the various archive’s crawlers. This matches the observation in Table 4 of a poor coverage of the Bitly URIs by the search engine indexes. The DMOZ sample has many URIs with more than 10 mementos. There are two reasons for this. First, DMOZ is a primary input for the Internet Archive’s crawler; and second, the DMOZ sample is retrieved from historical DMOZ archives which means these URIs may have existed a long time. Figure 3 shows the histogram for the number of mementos per URI-R.


Figure 4 shows the density of mementos per datetime. The x-axis represents the estimated creation date of the URI. The y-axis represents the number of mementos for a URI-R. The figure is supplemented with three density guidelines for 0.5, 1 and 2 and mementos per month.

Table 6 shows the correlation between the number of mementos and the number of backlinks. Backlinks are the incoming links to any URI. The number of backlinks is one indication of the popularity or importance of a website or page. The number of backlinks was calculated using the Yahoo Boss APIs. The correlation was calculated using Kendall’s τ and showed a weak positive relationship between the number of backlinks, and the number of mementos.


Table 7 contains statistics about the retrieved mementos

Table 5: Memento’s aggregator proxies list.


| Archive name | URI | Description |
|--|---|--|
| Internet Archive | http://www.archive.org | Internet Archive is a non-profit that was founded on 1996 to build an Internet library. Its purposes include offering permanent access to historical collections that exist in digital format |
| Google | http://www.google.com | Google is a search engine provides a cached version. |
| Yahoo | http://www.yahoo.com | Yahoo is a search engine provides a cached version. |
| Bing | http://www.bing.com | Bing is a search engine provides a cached version. |
| Archive-It | http://archive-it.org | Archive-It is a subscription service that allows institutions to build and preserve collections of digital content |
| The National Archives | http://nationalarchives.gov.uk | The National Archives is the UK government’s official archive. |
| National Archives and Records Administration | http://www.archive.gov | National Archives and Records Administration is the record keeper of all documents and materials created during the course of business conducted by the United States Federal Government. |
| UK Web Archive | http://www.webarchive.org.uk | UK Web Archive contains UK websites that publish research, that reflect the diversity of lives, interests and activities throughout the UK, and demonstrate web innovation |
| Web Cite | http://www.webcitation.org | Web Cite is an on-demand archiving system for web-references, which can be used by authors, editors, and publishers of scholarly papers and books, to ensure that cited web-material will remain available to readers in the future. |
| ArchiefWeb | http://archiefweb.eu | ArchiefWeb is a commercial subscription service that archives websites |
| California Digital Library | http://webarchives.cdlib.org | Focused cultural archives sponsored by the CDL. |
| Diigo | http://www.diigo.com | Diigo is a social bookmarking site that provides archiving services; subsumed Furl in 2009. |




(a) DMOZ



(b) Delicious




(c) Bitly




(d) Search Engines


Figure 1: URI-Ms per month.




(a) DMOZ



(b) Delicious




(c) Bitly




(d) Search Engines


Figure 2: URI-M distribution.




(a) DMOZ



(b) Delicious



(c) Bitly



(d) Search Engines

Figure 3: URI-Ms per URI-R.

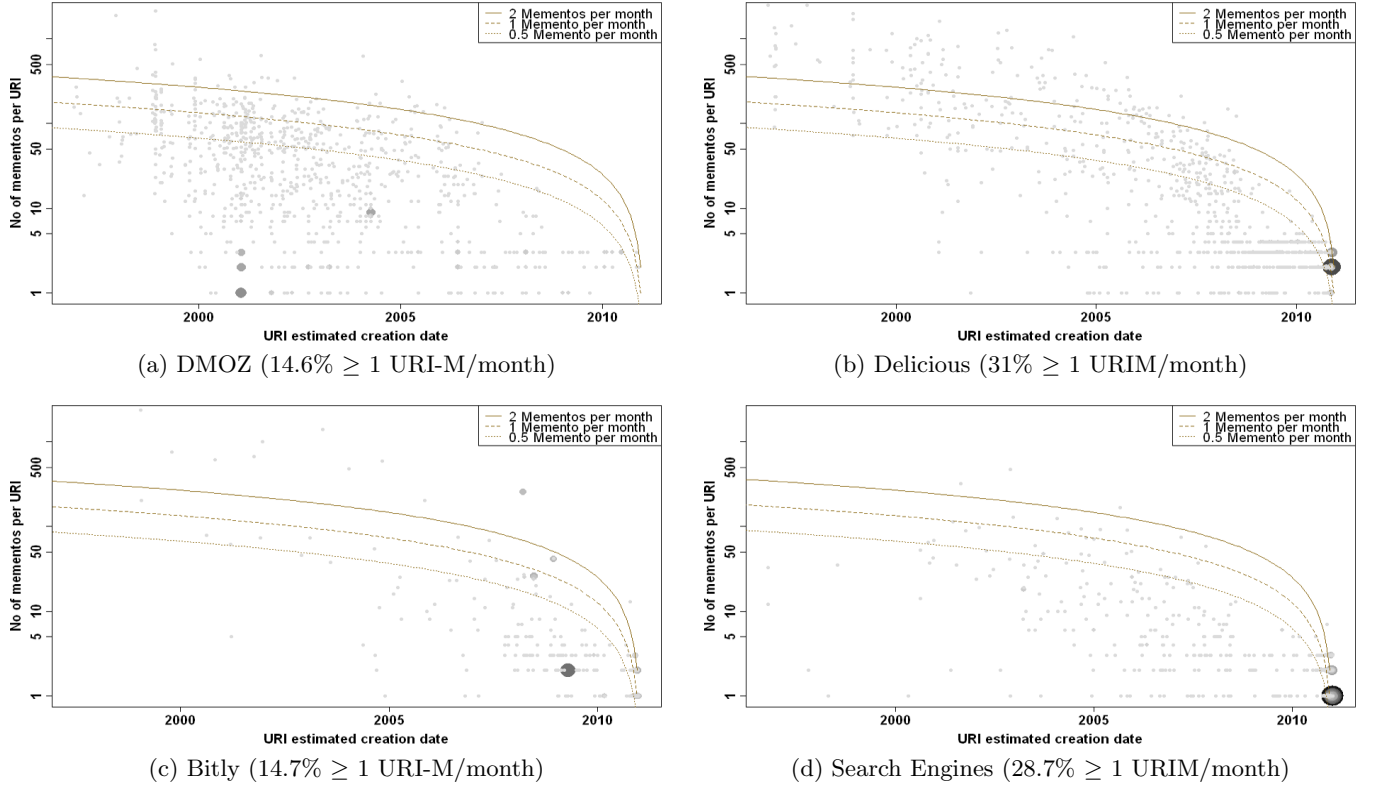


Figure 4: URI-M density to URI-R age (months).

Table 6: Mementos to backlinks correlation (τ).

| | DMOZ | Delicious | Bitly | SE |
|-------------|-------|-----------|-------|-------|
| Correlation | 0.389 | 0.311 | 0.631 | 0.249 |

by source. The table reports the number of URI-Ms from the source, the number of URI-Rs covered by the source, the mean URI-Ms per URI-R and corresponding standard deviation ($p \leq 0.01$).

5. ANALYSIS

To answer the question “How much of the Web is archived?”, the memento sources will be compared on 3 axis:

- *Coverage*. How many URI-Rs are archived?
- *Depth*. How many mementos are available?
- *Age*. How old are the mementos?

5.1 Internet Archive

The results showed that the Internet Archive (IA) has the best coverage and depth. Also, the IA covers URIs from 1996 to the present. IA has a delay from 6-24 months between the crawling and appearance on the archive web interface [9], but the results showed that this (6-24 months) period may be longer as less than 0.1% of the IA archived versions appeared after 2008. Moreover, IA provides archived versions for dead URIs (response 404).

In late 2010, Internet Archive launched a new WayBack Ma-

chine interface¹³. The new interface is supported with a new toolbar on archived pages and a new look of the calendar of page captures. The experiment results were retrieved using the classic WayBack Machine, because the new one is still in the beta version.

5.2 Search engine caches

Search engines (Google, Bing, and Yahoo) provide a cached version for most indexed pages [13]. Search engines provide a good coverage, but are limited to 1 copy per URI. Based on our observations, Google and Bing cached copies are kept for a maximum of one month. Yahoo provides cached version without date, we used a new technique to estimate the cached version age for Yahoo which may be several years [1].

5.3 Other archives

The other archives are mainly special purpose archives. These archives provide good coverage for their own web sites which may be limited to country (UK Web Archive, The National Archives, and NARA), special subscriber collections (Archive-It, ArchiefWeb, and CDLIB), or the user preferences (WebCite, and Diigo). Most of these sites have a high number of archived copies for these URIs. The age for the archived copies depends on the age of the archive itself.

5.4 Archive overlap

We also looked at the relationships and overlap between sources, which is shown in figure 5. The diameter of each cir-

¹³<http://waybackmachine.org>

Table 7: Mementos coverage per type.

| | DMOZ | | | | Delicious | | | |
|------------------------|--------|--------|-------|-----|-----------|--------|--------|-----|
| | #URI-M | #URI-R | Mean | SD | #URI-M | #URI-R | Mean | SD |
| Internet Archive | 55293 | 783 | 70.62 | 130 | 74809 | 408 | 183.36 | 325 |
| Google | 523 | 523 | 1 | 0 | 897 | 897 | 1 | 0 |
| Bing | 427 | 427 | 1 | 0 | 786 | 786 | 1 | 0 |
| Yahoo | 418 | 418 | 1 | 0 | 479 | 479 | 1 | 0 |
| Diigo | 36 | 36 | 1 | 0 | 354 | 354 | 1 | 0 |
| Archive-It | 92 | 4 | 23 | 41 | 500 | 38 | 13.16 | 30 |
| National Archives (UK) | 25 | 8 | 3.125 | 3 | 521 | 102 | 5.11 | 10 |
| NARA | 5 | 5 | 1 | 0 | 31 | 19 | 1.63 | 1 |
| UK Web Archive | 8 | 5 | 1.6 | 1 | 391 | 38 | 10.29 | 16 |
| Web Cite | 26 | 5 | 5.2 | 8 | 594 | 57 | 10.42 | 49 |
| ArchiefWeb | - | - | - | - | 22 | 3 | 7.33 | 11 |
| CDLIB | - | - | - | - | 20 | 5 | 4 | 4 |

| | Bitly | | | | Search Engines | | | |
|------------------------|--------|--------|--------|-----|----------------|--------|-------|----|
| | #URI-M | #URI-R | Mean | SD | #URI-M | #URI-R | Mean | SD |
| Internet Archive | 8947 | 70 | 127.81 | 406 | 4067 | 170 | 23.92 | 49 |
| Google | 253 | 253 | 1 | 0 | 486 | 486 | 1 | 0 |
| Bing | 204 | 204 | 1 | 0 | 515 | 515 | 1 | 0 |
| Yahoo | 87 | 87 | 1 | 0 | 229 | 229 | 1 | 0 |
| Diigo | 61 | 61 | 1 | 0 | 10 | 10 | 1 | 0 |
| Archive-It | 75 | 13 | 5.77 | 8 | 49 | 12 | 4 | 5 |
| National Archives (UK) | 531 | 12 | 44.25 | 145 | 1 | 1 | 1 | 0 |
| NARA | 10 | 2 | 5 | 6 | 4 | 2 | 2 | 0 |
| UK Web Archive | 2892 | 32 | 90.38 | 187 | 9 | 3 | 3 | 3 |
| Web Cite | 989 | 58 | 17.05 | 82 | - | - | - | - |
| ArchiefWeb | 609 | 1 | 609 | 0 | - | - | - | - |
| CDLIB | - | - | - | - | - | - | - | - |

cle represents relative coverage. The color density expresses the depth. The width of the edges expresses the intersection between both archives. The visualization was done by MapEquation¹⁴.

6. CONCLUSION AND FUTURE WORK

When we began the research to answer the question “How much of the Web is archived?” we did not anticipate a quick, easy answer. Suspecting that different URI sample sources would yield different results, we chose sources we believed would be independent from each other and provide enough insight to allow estimation of a reasonable answer. Instead we have found that the URI source is a significant driver of its archival status.

Consider the graphs in Figure 2. Clearly, the archival rate (URIs with at least one memento) is much higher for the DMOZ and Delicious samples than for the Bitly and search engine samples, which also differ considerably from each other. URIs from DMOZ and Delicious have a very high probability of being archived at least once. On the other hand, URIs from search engine sampling have about 2/3 chance of being archived and Bitly URIs just under 1/3. This leads us to consider the reasons for these differences.

Something that DMOZ and Delicious have in common is that a person actively submits each URI¹⁵ to the service. (Indeed DMOZ and Delicious can be both considered direc-

tory services, with DMOZ representing Web 1.0 and Delicious representing Web 2.0.) Search engine sample URIs, however, are not submitted in the same way. The process used by search engines is more passive, with URI discovery depending on search engine crawl heuristics. Bitly is more of a mystery. Our research did not delve into how Bitly URIs are used. But the low archival rate leads us to think that many private, stand-alone, or temporary resources are represented. These substantial archival rate differences have led us to think that publicity interest a URI receives is a key driver of archival rate for publicly-accessible URIs.

Future work will include study of the relationship between the rate of change of the web page and the rate of the archiving process. Also, we plan to study of the quality of the archive itself. This work has been done on a general sample of URIs. In future work, the archived URIs will be studied based on specific languages beyond English.

7. ACKNOWLEDGMENTS

This work is supported in part by the Library of Congress. We would like to thank Herbert Van de Sompel and Robert Sanderson from Los Alamos National Laboratory, and Kris Carpenter Negulescu and Bradley Tofel from Internet Archive for their positive comments and explanations.

8. REFERENCES

- [1] Ahmed AlSum and Michael L. Nelson. Validating Yahoo’s API last modification date, 2011.
- [2] Ziv Bar-Yossef and Maxim Gurevich. Random sampling from a search engine’s index. *Journal of the*

¹⁴<http://www.mapequation.org>

¹⁵Spam-bots notwithstanding

- ACM (JACM)*, 55(5), 2008.
- [3] C. Casey. The Cyberarchive: a Look at the Storage and Preservation of Web Sites. *College & Research Libraries*, 59:304–310, 1998.
 - [4] Michael Day. Preserving the fabric of our lives: A survey of web preservation initiatives. In *Proceedings of the 7th European Conference on Research and Advanced Technology for Digital Libraries, ECDL 2003*, pages 461–471, 2003.
 - [5] Alex Franz and Thorsten Brants. All our n-gram are belong to you. <http://googleresearch.blogspot.com/2006/08/all-our-n-gram-are-belong-to-you.html>.
 - [6] Daniel Gomes, Sérgio Freitas, and Mário Silva. Design and selection criteria for a national web archive. In Julio Gonzalo, Costantino Thanos, M. Verdejo, and Rafael Carrasco, editors, *Research and Advanced Technology for Digital Libraries*, volume 4172 of *Lecture Notes in Computer Science*, pages 196–207. Springer Berlin / Heidelberg, 2006.
 - [7] A. Gulli and A. Signorini. The indexable web is more than 11.5 billion pages. In *Special interest tracks and posters of the 14th international conference on World Wide Web, WWW '05*, pages 902–903, New York, NY, USA, 2005. ACM.
 - [8] Internet Archive. “how can i get my site included in the wayback machine?”. <http://www.archive.org/about/faqs.php#1>. Accessed January 17, 2011.
 - [9] Internet Archive. “Why are there no recent archives in the Wayback Machine?”. <http://www.archive.org/about/faqs.php#103>. Accessed January 17, 2011.
 - [10] Julien Masanès. *Web Archiving*. Springer, 2006.
 - [11] Ingrid Mason. Virtual preservation: How has digital culture influenced our ideas about permanence? changing practice in a national legal deposit library. *Library Trends*, 56(1):198–215, Summer 2007.
 - [12] Frank McCown and Michael L. Nelson. Agreeing to disagree: search engines and their public interfaces. In *Proceedings of the 7th ACM/IEEE-CS joint conference on Digital libraries, JCDL '07*, pages 309–318, New York, NY, USA, 2007. ACM.
 - [13] Frank McCown and Michael L. Nelson. Characterization of search engine caches. In *Proceedings of IS&T Archiving 2007*, pages 48–52, May 2007. arXiv:cs/0703083v2.
 - [14] Frank McCown and Michael L. Nelson. A framework for describing web repositories. In *Proceedings of the 9th ACM/IEEE-CS joint conference on Digital libraries, JCDL '09*, pages 341–344, New York, NY, USA, 2009. ACM.
 - [15] Frank McCown, Joan A. Smith, Michael L. Nelson, and Johan Bollen. Lazy preservation: Reconstructing websites by crawling the crawlers. In *WIDM '06: Proceedings of the 8th annual ACM international workshop on Web information and data management*, pages 67 – 74, 2006.
 - [16] G. Monroe, J. French, and A. Powell. Obtaining language models of web collections using query-based sampling techniques. *Hawaii International Conference on System Sciences*, 3:67b, 2002.
 - [17] Kris Carpenter Negulescu. Web archiving @ the Internet Archive. http://www.digitalpreservation.gov/news/events/ndiipp_meetings/ndiipp10/docs/July21/session09/NDIIPP072110FinalIA.ppt, 2010.
 - [18] C. Olston and S. Pandey. Recrawl scheduling based on information longevity. In *Proceeding of the 17th international conference on World Wide Web*, pages 437–446. ACM, 2008.
 - [19] Mike Thelwall and Liwen Vaughan. A fair history of the web? examining country balance in the internet archive. *Library & Information Science Research*, 26(2):162–176, 2004.
 - [20] Herbert Van de Sompel, Michael Nelson, and Robert Sanderson. Http framework for time-based access to resource states – memento, November 2010. <http://datatracker.ietf.org/doc/draft-vandesompel-memento/>.
 - [21] Herbert Van de Sompel, Michael L. Nelson, Robert Sanderson, Lyudmila L. Balakireva, Scott Ainsworth, and Harihar Shankar. Memento: Time travel for the web. Technical Report arXiv:0911.1112, 2009.
 - [22] Herbert Van de Sompel, Robert Sanderson, Michael Nelson, Lyudmila Balakireva, Harihar Shankar, and Scott Ainsworth. An HTTP-based versioning mechanism for linked data. In *Proceedings of Linked Data on the Web Workshop (LDOW2010)*, April 27 2010. http://events.linkedata.org/ldow2010/papers/ldow2010_paper13.pdf.




Figure 5: Archive connections graph